Evaluation of Supervisory vs. Peer-Peer Interaction for Human-Robot Teams

Julie L. Marble,
Human, Robotics and Remote Systems Group
Idaho National Engineering and Environmental Laboratory
Idaho Falls, ID 83415
208 526 1585
marbjl@inel.gov

David J. Bruemmer
Human, Robotics and Remote Systems Group
Idaho National Engineering and Environmental Laboratory
Idaho Falls, ID 83415
208 526 4078
bruedj@inel.gov

Douglas A. Few,
Human, Robotics and Remote Systems Group
Idaho National Engineering and Environmental Laboratory
Idaho Falls, ID 83415
208 526 3077
fewda@inel.gov

Donald D. Dudenhoeffer
Human, Robotics and Remote Systems Group
Idaho National Engineering and Environmental Laboratory
Idaho Falls, ID 83415
208 526 0700
E-mail: dudedd@inel.gov

Evaluation of Supervisory vs. Peer-Peer Interaction with Human-Robot Teams

Abstract

We submit that the most interesting and fruitful human robot interaction may be possible when the robot is able to interact with the human as a team member, rather than as tool. To that end, the INEEL has developed a dynamic autonomy control architecture, which allows the user to select the level of robot initiative. However, the benefits of shared control can all too easily be overshadowed by control challenges inherent to blending human and robot initiative. The most important requirement for peer-peer interaction is trust. For operators to embrace initiative taking autonomous robots, controls must be crafted to operate and fail predictably. The human must be able to predict the robot's behavior and understand the reason and effects of robot initiative. This paper discusses the findings of recent human participant usability testing which takes our current implementation to task using a pseudo-search and rescue scenario within a complex, real-world environment. The purpose of this testing was to examine how the human operator prefers to work with the robotic system at each level of autonomy, and how interaction with the robot should be accomplished to best support situation awareness and task completion. Both novice and experts in remote system deployment were used in the assessment of the different control modes. Critically, analyses revealed that the architecture equally supported situation awareness and target detection by novices and experts, although experienced users were more likely to have heightened performance expectations of the interface. Through this human-participant testing, the INEEL explored the potential for humans and robots to blend their initiative. This paper goes on to report several surprising results regarding the ability of participants to effectively utilize the collaborative workspace and, most importantly, their ability to understand and their willingness to accept robot initiative.

I. Introduction

A nationwide attempt to accelerate cleanup efforts at DOE sites has given rise to an unprecedented need to remotely characterize buildings that have been marked for decontamination and decommissioning. Although, teleoperated robotic solutions offer significant benefits in terms of human exposure, time, cost and quality of data, close evaluation of these operations brings to light severe limitations to the master-slave strategy employed, including lapses in communication and situation awareness, which result in damage to the environment, loss of or damage to the robot, requiring personnel to enter the environment [1], [2]. As mechanical 'subordinates,' such teleoperated robots are dependent on continuous, low-level input from a human and are poorly equipped to cope with communication failures or changes in operator workload.

On the other hand, attempts to build and use autonomous systems have failed to acknowledge the inevitable boundaries to what the robot can perceive, understand, and decide apart from human input. Human Factors studies in the area of human-computer interaction (HCI) and human-machine interaction (HMI) reveal that complex tasks are more successfully performed when the system is designed to support the needs of the human rather than eliminating the human from the system [2]. In fact, in many cases, the goal to eliminate the human from the system has resulted in significant system failures including loss of life specifically because the system was not designed to support interaction with the human [3]. Within the field of robotics, both fully autonomous approaches and teleoperated approaches have failed to realize the immense potential for humans and robots to work together as a team where each member is invested with agency -- the ability to actively and authoritatively take initiative accomplish task objectives. Such *mixed-initiative* systems can support a spectrum of control levels, allowing the human and robot to support each other in different ways as needs and capabilities change throughout a task. To accomplish this objective, mixed-initiative robots should:

- Possess intrinsic intelligence and agency;
- Protect humans, environment, and self;
- Dynamically shift the level of autonomy;
- Accept different modes of human intervention; and
- Recognize when help is needed.

Towards these aims, research efforts at the INEEL have developed a novel robotic system that can leverage its own intelligence to support a spectrum of control levels. Rather than conceive of machines as passive tools or, conversely, as totally autonomous entities that act without human intervention, we believe that it is more effective to consider the machine as part of a dynamic human-machine team. The unanswered question is whether and how such control schemes can be implemented effectively. This paper discusses a human-participant usability experiment where human operators teamed with a robot in order to find primary and secondary search targets within a remote building. The goal was to examine the benefits and pitfalls to the different modes of operator intervention with special attention to the question of whether operators are more successful with supervisory control or peer-peer interaction and which they prefer.

II. Mixed-Initiative Teaming

Given the desire to employ robots in hazardous, high-risk environments, the ability for robot and human team members to work as peers presents particularly difficult HRI questions:

- Should a machine give orders to humans?
- What needs to be communicated and at what level of abstraction?
- Can team roles shift to reflect changing capabilities?
- When should machines say no to humans or to other machines?

Often the robot, placed within the environment is better equipped to make decisions regarding navigation. This is especially true when thick concrete shielding prevents high-bandwidth communications such as video from reaching the human user in a timely manner. The human user may rely on the self-protective, autonomous capabilities of the robot while giving commands based on less detailed information, such as a map of the environment created on the fly. Yet, at times, the human team member must override these self-protective capabilities to accomplish a critical mission. For instance, when faced with a an obstacle on the path, the user may shift the robot out of the leadership responsibility for navigation in order to attempt to move to object out of the way, but the robot can retain the “right” to refuse human commands when its own safety or that of the environment is at stake. The robot uses an assessment of the physical resistance to motion as well as bump sensors, short-range infrared break beams, tilt sensors, electrical current draw, and other means to monitor and protect itself. Throughout a task, the roles of each team member are bounded by a complex and changing web of capabilities and limitations to which each member must adapt and respond. However, to support effective team performance, this information must be available to all members, when it is needed. The critical issue then is to determine when is it needed.

Within a healthy team, group roles evolve dynamically to meet new and unforeseen challenges. When we apply team role theory to human-robot interaction, we find that very few true human – robot teams exist. Machines are generally used as tools, rarely as peers and scarcely ever as task leaders [4]. Just as roles evolve between human team members, mixed human – robot teams should be able to modulate the level of robotic initiative in order to balance changes in the environment, task, and capabilities of other team members.

The need to transition roles and responsibilities among the members of mixed human-robot team presents new challenges. After all, the dynamics of performance for robots are drastically different than for humans. Performance of a robot may far exceed human abilities in certain elements of a task; however, robots are notoriously unable to degrade gracefully in the face of component failure or unforeseen changes in the environment or task. Thus, while robots may be more reliable than humans within certain parameters, they often lack the reliability necessary for tasks where such parameters cannot be guaranteed. The answer to this challenge is not to keep robots as passive tools; rather, we must have a team dynamic where roles can shift to compensate for the unique kinds of failures and limitations possessed by each member.

The need for both human and robot to predict and understand one another's actions presents a daunting challenge. For each level of robot initiative, the user must be able to quickly and accurately predict robotic responses and understand how cumulative robotic actions may converge to fulfill task objectives. Just as the human develops an understanding of the robot's behavior, the robot must be able to understand and predict the behavior of human members of the team. The robot's expectations must allow it to recognize human limitations and anticipate human needs without second-guessing the human's every move. When robots do intervene, the human's understanding of robot behavior must be able to explain why the robot has stepped in and what this shift in control means for the task at hand. This need for dynamic role switching makes critical the consideration of human factors in the design of the human-machine interface.

III. Experimental Design

A. Participants

Participants were recruited from the available pool of INEEL site employees. Eleven INEEL employees ranging in age from 32 to 57, with median age 41, participated. All participants were paid their normal salaries for their time, which was not more than 2 hours. Participants included one female and ten males. Seven participants had no prior experience with remote or robotic systems (never used any type of remote system) or some experience (had used a remote system such as a remote radio frequency-controlled car as a hobby). Four expert users defined by extensive (average 7 years) primarily job related experience using teleoperated vehicles, master-slave systems, or similar systems also participated.

B. Test Location

Testing took place at Chemical Processing Plant (CPP) 1662 of the INEEL, which houses the Remote systems branch of the Human, Robotic, and Remote systems group. The layout of the building is shown schematically in Figure 1, below. The building is made of steel and concrete, and which does disrupt video feed signals, as could often be encountered in a "real world setting." The building is approximately 3600 square feet, 60 feet long by 60 feet wide. The central hallway is approximately 5 feet wide and 40 feet long; the distance from the left turn to the building exit is approximately 20 feet. This turn was narrowed by coveralls and boots placed in a hanging area. Opposite these, a box (18" x 12") was placed next to the doorway, depicted in the schematic, to further narrow the hallway. The high bay area from which the operator controlled the robot is approximately 30'w x 60' l. The main storage room in which 2 primary targets were located is approximately 20 x 20 feet. Just outside the storage room doors, the hallway was narrowed by two trash cans placed 4.25 feet diagonally from one another, as shown in the schematic. It was possible for the operator to use the robot to push these trash cans out of the way, however, they were asked to slalom between them to move down the hallway. The third primary target was located in the janitorial room at the end of the central hallway. Opposite this doorway, a fire extinguisher was hung on the wall, which may have been low enough to be detected by the robot's sensors. Although the robot had to move into the storage room to identify the first 2 primary targets, it was not necessary for the robot to enter the janitorial closet to identify the third target. Doors to the two offices were typically closed, although because of worker activity in the building, they were occasionally open. Operators were asked to not allow the robot to enter the offices.

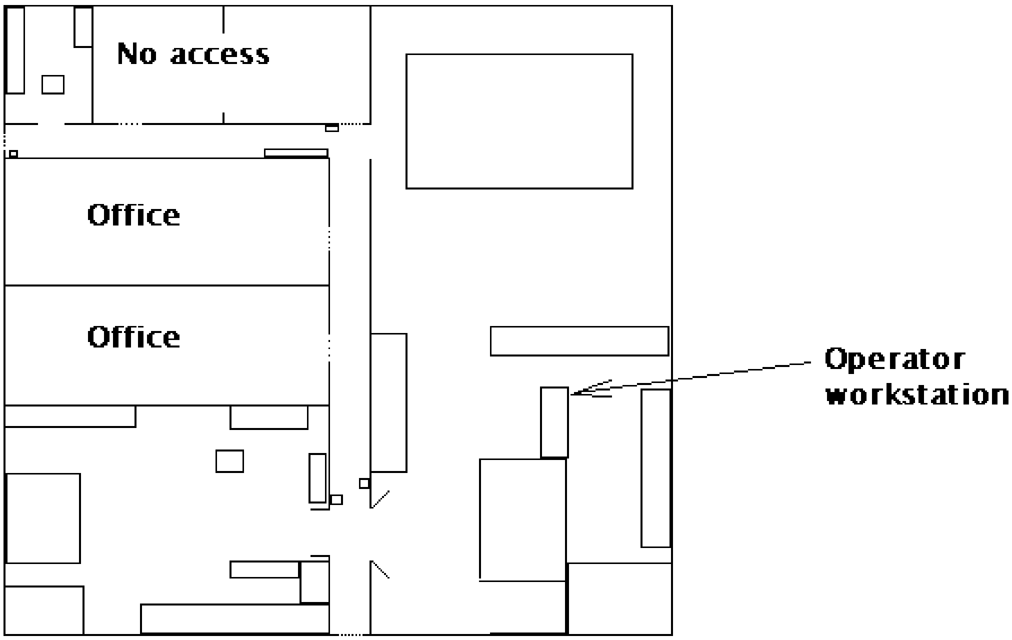

Figure 1. Schematic layout of CPP 1662. (Not to scale.)

C. System Design

Figure 2. The ATRV Jr. and component sensors

The robot used in these tests was an all terrain remote vehicle ATRVJr (shown in Figure 2), with an architectural control platform developed at the INEEL. The robot is equipped with a variety of range sensor information including inertial sensors, compass, wheel encoders, laser range finders, computer vision, infrared break beams, tilt sensors, bump sensors, and sonar. Streaming video from the robot is provided to the operator. In addition, the sensory information noted above is fused to provide a map of the area created on the fly as the robot moves through the environment (shown in the upper right quadrant of Figure 3), the map is oriented so that up is north. The robot continuously evaluates this information and provides information on the state of the sensors to the user. The robot abstracts information about the environment at many levels including terse textual descriptions of the robot's local surroundings and the choices (depending on the level of autonomy) that face the human user.

The architecture controlling the robot supports four levels of increasing robotic autonomy, the first three of which were used in the usability test:

> ***Teleoperation***: The user has full, continuous control of the robot at a low level. The robot takes no initiative except to stop once it recognizes that communications have failed. It does indicate the detection of obstacles in its path to the user, but will not prevent collision.
>
> ***Safe Mode****:* User directs movements of robot, but the robot takes initiative and has the authority to protect itself based on its proprioception and self-status evaluation; for example, it will stop before it collides with an obstacle, which it detects via multiple sensors. The robot will refuse to undertake a task if it cannot safely accomplish it.
>
> ***Shared Control****:* The robot takes the initiative to choose its own path in response to general direction input from the operator. Although the robot handles the low level navigation and obstacle avoidance, the user supplies intermittent input, often at the robot's request, to guide the robot in general directions.
>
> ***Full Autonomy:*** The robot performs global path planning to select its own routes, acquiring no operator input except high-level tasking such as "follow that target" (specified using a bounding box within the visual display) or "search this area" (specified by drawing an area within the map interface module). If the operator uses the pursuit button, the robot will autonomously follow whatever object the user identifies within the visual image. (Although this mode was available, it was not used during the test.)

A fourth mode of interaction used in this experiment was ***Dynamic Control***. This is not a level of autonomy in itself. Rather it refers to the ability of the user to switch between any of the above modes of autonomy to best accomplish the task by selecting the mode from the touch screen. For the usability task, participants were asked not to utilize the fully autonomous mode.

For each level of autonomy, perceptual data is fused into a specialized interface presented on a touch screen (shown in Figure 3) that provides the user with abstracted graphical and textual representations of the environment and task appropriate for the current mode. Immediate obstacles that inhibit motion are shown as red ovals to the side or to the front or back and resistance to motion is shown with arcs emanating from the wheels, shown for the rear wheels on the graphic depiction of the robot in the lower right quadrant of Figure 3. Video information is provided to the operator continuously. As the operator touches the visual image on the display, the robot's camera aligns to center that part of the image. The operator can also manipulate the camera by selecting the tilt and pan buttons located around the video display. The robot relays synthesized, high-level information (such as suggestions and requests for help) to the user in a textual form using the feedback textbox below the image window.

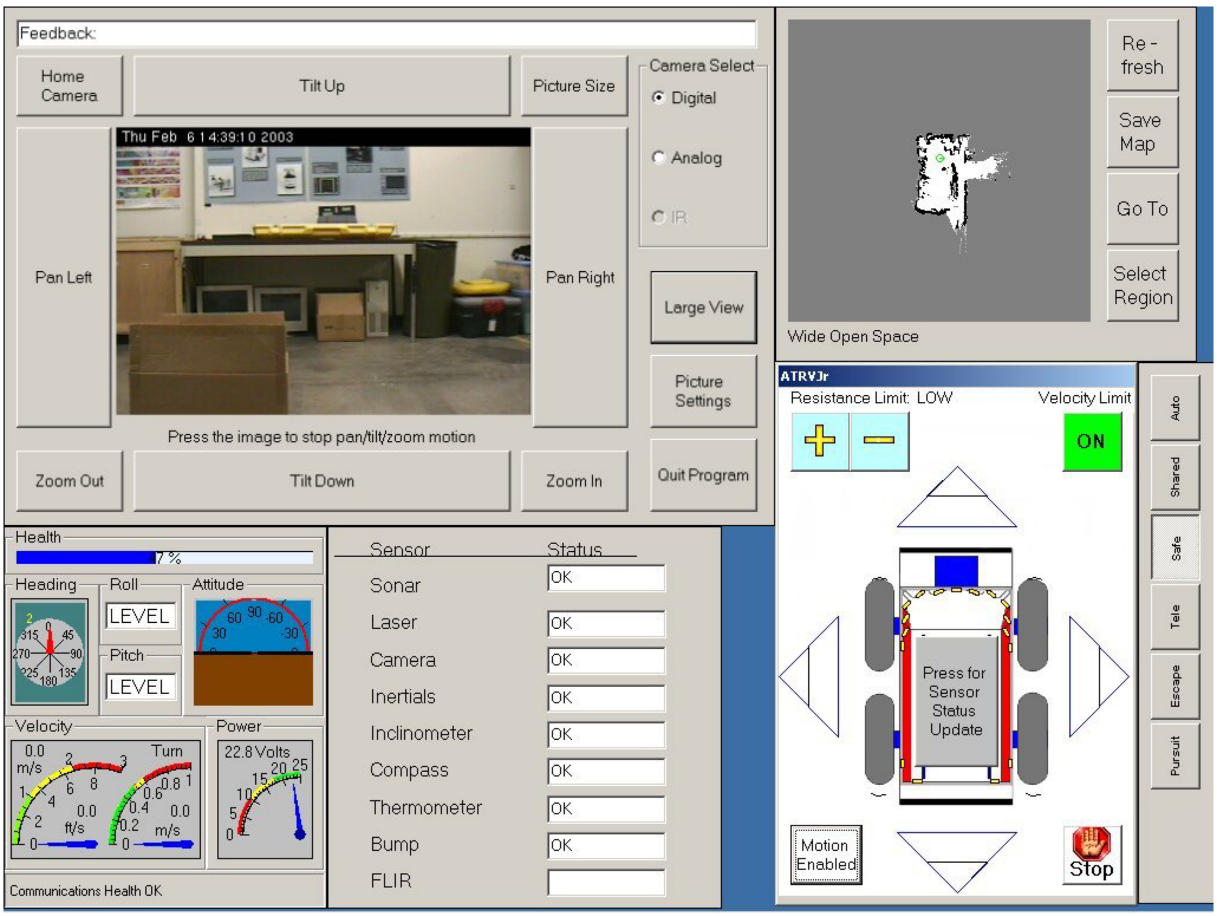

Figure 3. Current human-robot interface.

Also note that the robot provides textual reports on environmental features at the bottom of the map window and reports on communications status at the bottom of the robot status window. The robot status window (lower left of Figure 3) provides a variety of information including pitch and roll, power, heading, speed and a fusion of this information into a single measurement of "health." The user may direct the robot by touching the arrows or using a joystick, both of which were available during the test.

D. Method

Participants were asked to search a building using a robot equipped with multiple levels of autonomy to locate 3 targets in 3 pre-specified locations as quickly and safely as possible. Targets were selected at random from a set of 20 different stuffed animals varying in size from 6 inches tall to approximately 1 foot tall. (Stuffed animals were chosen because of their varying size and because they would obviously not be part of the office environment being searched.) Participants made four searches of the building; in three searches, participants were limited to only one level of autonomy (teleoperation, safe teleoperation, and shared autonomy) and the order of use was randomized, while in the fourth search, participants were allowed to shift the level of autonomy as much as they desired to accomplish the task (dynamic autonomy).

During each search, two alternate targets were placed along the route the robot would need to navigate, although the locations of these secondary targets were randomized on each trial. Participants were told to be aware that up to two of these secondary targets might be present, but that they were not required to find them. These targets were used to assess the degree of situation awareness participants had about the environment of the robot while engaged in navigation.

A secondary math task was used to assess how taxing the use of each search mode was on the participants. During each search, participants were also asked to solve simple (one digit by one digit) multiplication problems. The problems were presented on a laptop computer placed next to the touch screen display. Problems began appearing 15 seconds into the search, and 5 seconds after the previous problem was answered. As each problem was presented, a quiet tone sounded to alert the operator. The program recorded the number of problems answered by the operator, the number correct, and the time from presentation to answering. Such an intermittent secondary task is similar to the types of interruptions an operator might encounter during 'real-world' use of the system during a search, perhaps by questions from a team member. Participants were instructed that the math task was not as important as performing the search task safely and quickly, and that if the search task took all of their attention to disregard the math problems until the search task was less difficult.

During performance of the search task, the joystick, touch screen, and laptop were videotaped to record the actions of the operator. In addition, a trained observer remained with the participant to note any comments that may not have been recorded by the video. After each search, the participant was asked to fill out a 19 item subjective assessment regarding the search task just completed. This can be seen in Appendix A. Finally, a cameraperson followed the robot throughout the building to correlate the actions taken by the participant with the actions of the robot. The cameraperson was to intervene and stop the robot only if damage to the building walls or a person were imminent. If the cameraperson had to intervene, he hit the emergency brake on the robot, came to tell the participant, then moved the robot into a position that was no longer in contact with the walls. This added a 'penalty' of approximately 20 seconds onto the time for each participant. During instruction, the time to perform the searches was emphasized. The number of interventions per participant and mode was recorded.

Before beginning the searches, participants were given 20 minutes to familiarize themselves with the interface and behavior of the robot in the high bay area of the building where the participant workstation was located. During this familiarization period, the different aspects of the display were explained, as well as the different levels of autonomy. Participants were then asked to perform various tasks using each level of autonomy (e.g., attempt to move a trash can, which was only possible in the teleoperated mode), until the participant indicated that they felt that they knew how to use the system. Although this time is not sufficient to completely train an individual to use the system, it is realistic because an end user would most likely be trained once on the system, but then only have reason to use the system every few months at most for search and rescue tasks.

IV. Results

A. Observational and qualitative results.

One interesting result obtained from the testing was how participants differed in their approaches to autonomy. Generally participants could be divided along one line: those who were willing to release the controls after giving a general direction command in the shared autonomy mode, and those who would continue to give direction commands in shared mode as in the teleoperated mode. Participants who were more successful in the shared mode tended to be those who appeared more willing to actually release the controls. A second factor noted was that participants experienced with tethered robotic systems often spent the first few minutes with the robot adjusting the camera down and panning around the floor. Interviews with these participants revealed that they were attempting to learn how much of the body of the robot could be seen with the camera, in order to determine how tight an area they could maneuver within. All participants indicated a desire for the video display to have a depth indicator on it, especially in teleoperated and safe modes, the implementation of which is being investigated. A final grouping factor was that some participants quickly learned to navigate around obstacles in both teleoperated and safe modes, not based on the video information, which was occasionally interrupted, but rather based on the obstacle indicator on the interface to the almost complete neglect of the video information. Several participants were highly successful and very quick in navigating corners and the slalom based entirely on this information, while others did not appear to notice its utility or importance. Participants who did navigate around obstacles in this manner appeared to neglect the video until they had reached space in which none of the obstacle indicators were lit, typically once they had completed going around the objects. It is not clear at this time what separates these two groups. Finally, it should be noted that in teleoperated mode, no participants thought to have the robot move obstacles out of the path, although they were asked to purposeful do this during the familiarization session. This may have been caused by being told that they would need to navigate around obstacles during the walk-thru of the building, but it is interesting that no participants asked whether they could move the obstacles in the slalom.

Also intriguing was the observation that the oldest participant adapted most quickly to the highest level of autonomy, and had the greatest difficulty with the teleoperated mode. Interviews with this participant revealed that he has over 17 years of experience with hydraulic cranes, and to some extent was expecting similar force feedback from the controls in teleoperated mode and safe mode. Based on this finding, we are examining whether to introduce a form of force feedback into the controls. In the shared autonomy session, this participant would frequently point the robot in the general direction he wished it to move, and then physically release the controls to devote his attention to the math task for up to 15 seconds. During this time, the robot would navigate in the desired direction, which in some cases meant backing up slightly to find more open space. This type of reaction was frustrating for some participants, frequently those who had the most experience with tethered vehicles, who would fight the controls to force the robot to continue in the desired direction, which was difficult because the robot would not perceive enough space to continue solely in that direction.

B. Quantitative results.

Feeling of being in control of the robot increased across sessions, $F(2, 10) = 8.613$, $MSE = 4.58$, $p < 0.001$, with the feeling of control being least in session 1 and increasing regularly to session 4, indicative of a typical learning curve. There was a moderate interaction of session and mode for feeling in control, $F(2, 10) = 2.637$, $MSE = 1.402$, $p = 0.055$, indicating that participants felt in greatest control on the first session when they used the safe mode. Intriguingly, there was no difference in feeling of being in control between teleoperated mode and shared mode on the first session, although there was a significant difference in the ratings by the second session. There was no difference in ratings for the safe and teleoperated modes in the third session, and feeling of control increases for the shared mode when used in the third session. This implies that learning about the system increases feelings of control of the system, regardless of the mode of use. However, there was a small but not significant decrease in ratings of control for the fourth dynamic mode. This may have been due to the user needing to select the mode s/he was most comfortable with, which increased by a small amount the attention the user must give to the task.

Feeling of control varied across modes of control, $F(2, 10) = 12.288$, $MSE = 6.534$, $p < .0001$, with the greatest feeling of control occurring in dynamic mode (which was always the final session), next safe mode, tele mode next, and shared mode yielding the least feeling of control. That the shared mode yielded the least feeling of control is unsurprising because in this mode the robot had the most control of navigation while the operator gave general direction commands, rather than controlling the exact path of the robot as in teleoperated and safe modes. Interesting to note, ratings of feelings of control were identical during session 1 for both the teleoperated and shared modes, while safe mode yielded the greatest ratings of control during session 1.

Participants were able to locate and identify all primary targets, except in 4 cases; one where communications interference prevented completion of the task, two in which the participant chose to end the task because they were unable to navigate the slalom of the hallway, and one in which the participant apparently forgot to seek out target 2. Analysis of mean certainty ratings of identifying the secondary targets were analyzed with respect to session and autonomy mode revealed no differences, however, because of the small sample size, this may be due to lack of statistical power. Participants tended to be more certain of having identified the secondary targets in session 2 and in safe mode than in other sessions or modes, although this was not statistically significant. Mean certainty ratings of spotting the secondary targets were then analyzed with respect to the actual number of secondary targets identified.

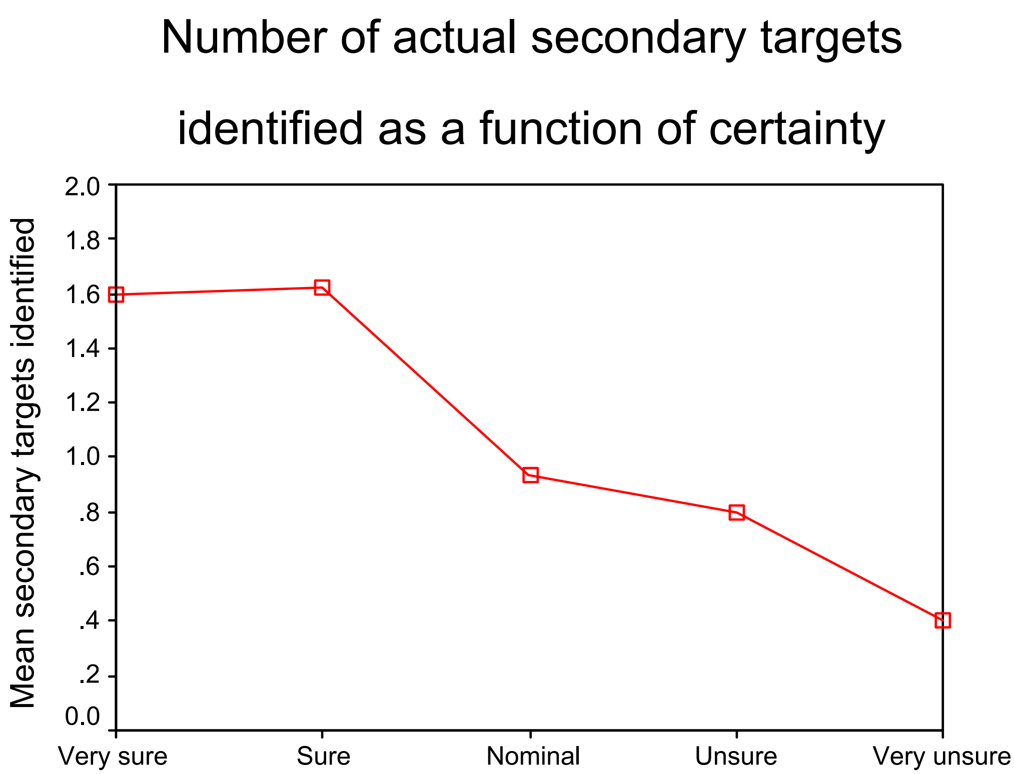

Certainty of having identified all the secondary targets was a significant, positive indicator of actual performance finding the secondary targets, $F(4, 31) = 3.469$, $MSE = 1.640$, $p < 0.022$. In other words, the more certain participants were that they had identified all secondary targets then the more likely they were to have actually identified all secondary targets. This implies that the information presented in the interface was sufficient to allow the participants to maintain situation awareness, although the task itself was structured to emphasize speed of performance rather than certainty of completion. In addition, there was no effect of previous robotic experience on actual secondary targets identified. This further supports our contention that the interface and control structure is equally usable by all groups. This is a critical conclusion because in a real world situation, the users of such a robotic architecture (i.e., search and rescue personnel) will not be constantly retraining on the use of the system. Rather, such users will be trained once, presumably in a short time period, regarding how to use the system, and afterward will only put the system into actual use once every six months to a year. Therefore, it is critical that the system equally maintain the situation awareness of experts and novices in a search task.

Analysis of the ratings of the usefulness of the map indicated low utility regardless of mode and session. This was assumed to occur for three reasons. First, the building had an extremely simple layout, an L-shaped corridor with one large room to search at the beginning. Second, participants walked the building at the start of the test to be shown the locations of the targets. Third, participants were asked to keep the map in its default size, rather than enlarge it to its more useful maximal size that would allow easier identification of the direction of movement of the robot. This constraint was made because of concerns that participants would leave the map at maximum size during the test, which would inhibit their use of other aspects of the HRI. In addition to rating the utility of the map after each session, participants were interviewed regarding their usage of the map and video at the end of all trials. In future tests, it may be useful to perform eye tracking to determine the amount of time devoted to each segment of the screen.

Participants were also asked to rate the usefulness of the video feed for performance of the tasks. Because of the amount of steel in the building video communications were easily disrupted, as could occur in a real world task. This was intriguing, because it allowed us to assess the degree to which participants required video information in each mode of autonomy. Analysis of the subjective assessments revealed a main effect of skill, $F(1, 39)= 32.331$, $MSE = 17.962$, $p< 0.001$. On average, the more expert users indicated less utility of the video than did the novices. This may have been confounded in part because at the time of the usability test, the expert users contacted had been contracted to build a tethered system with high fidelity video and high resolution for deployment in a tank. During testing, several of the experienced users made comparative comments to that system.

This analysis also revealed a tendency toward a main effect of mode, $F(2, 39) = 2.903$, $MSE = 1.613$, $p = 0.074$, which indicated that participants felt the video feed was most useful in dynamic and safe modes, and the least useful when in the shared mode. However, there was also a tendency for an interaction of ratings of video utility between session and mode of autonomy, $F(4, 39) = 2.286$, $MSE = 1.27$, $p = 0.09$. Ratings of the utility of the video feed increased by session 3 for the shared mode, but decreased for teleoperated and safe modes. At session 2, ratings of utility were similar for the teleoperated and shared modes, but best for the safe mode; however by session 3 ratings of the video feed were worst for the safe mode, slightly better for the teleoperated mode, but best for the shared mode. This may have been due in part to a tendency especially noted in the novices to manipulate the camera by moving the robot, rather than manipulating the camera directly. In the shared mode, this was not possible as the robot continuously moved toward open space. Therefore, in order to manipulate the camera, the operators had to stop the robot and then manipulate the camera in order to survey the surroundings, which may have led to greater utility of the video feed. The interaction would have been caused in part by the previous requirement and also due to participants learning across sessions that it is easier to manipulate the camera directly than to position the robot for viewing items.

V. Conclusions

The initial analyses of the results of subjective and objective measures in the human-robot usability test have revealed a number of potentially important findings. First, participants maintained situation awareness in all modes of autonomy, regardless of experience with robotic systems. This was demonstrated by the strong positive relationship between the participants' certainty that they had found all secondary targets and the number of secondary targets they had identified. That situation awareness was maintained regardless of experience is critical because, as discussed previously, although an operator may initially be trained in the use of the system, that level of training will likely not be maintained when it is not necessary to use the system, which may be weeks or even months of time.

Furthermore, operators differed in how they were able to work with higher levels of autonomy. Some users seemed quite willing to utilize the autonomous capabilities of the robot while others continuously fought with the robot for control, losing efficiency to "mode switching." Although analysis of task completion times and robot path analysis is not yet complete, these initial analyses imply that the most successful operators were able to utilize higher degrees of autonomy. Some users quickly adapted to the autonomous capabilities of the robot while others preferred the teleoperated and safe modes simply because they required less understanding or conformed more closely to their expectations. Future experiments will provide at least some level of training for the participants. Although reduction in training time has often been given as a reason for developing robot intelligence, it may well be that the higher the level of robot autonomy, the more training is needed by the human operator to effectively understand how and why the robot will behave the way it does. Interesting parallels might be drawn by looking at how K-9 police officers learn to work as team members with their dogs. Such training is continuous, both to reinforce the desired responses in the animal, but also to reinforce the commands in the human team member.

Most intriguing were the differences in the need for video information across participants. Some participants appeared to navigate entirely based on the video information, while others were able to take advantage of other information in the display, such as the obstacle indicator. Because the environment was so simple, it did not appear necessary for the participants to take advantage of the mapping capability of the robot and interface. Future research will be performed in more complex environments with more similar landscape features in order to examine how participants interact with the mapping capability.

Also critical to the application of autonomous robotics was the indications that performance in the shared mode benefited from practice in the teleoperated and safe modes, although this may have been a function of learning about the task, rather than learning how to interact with the system. However, performance in the shared mode did improve more significantly across sessions when compared to teleoperated and safe modes. Future research will utilize more complex environments (e.g., lowered lighting, more obstacles) to force the operator to rely more greatly on the proprioceptive capabilities of the robot, in order to reveal more information on how people should interact with autonomous, initiative taking robotic systems. Although the usability study has illuminated many opportunities for improvement in experimental design and training domains, we believe that the collaborative, cognitive workspace offers a means to mediate between human and robot team members, provides a means to fuse sensing from differing modalities and communicates knowledge from disparate perspectives in support of remote search tasks.

References

[1] Anderson, M., Conner, C., Daniel, V., McKay, M., & Yancey, N., “Demonstration of the robotic gamma locating and isotopic identification device,” in *Proceedings from the American Nuclear Society Spectrum*, Reno, NV, August 2002.

[2] Abbott, K.A., Slotte, S.M., & Stimson, D.K., “Federal Aviation Administration Human Factors team report on: The interfaces between flightcrews and modern flight deck systems,” Federal Aviation Administration, Washington, DC, Technical Report, June 1996. [Online]. Available: http://www.faa.gov/avr/afs/interfac.pdf, pp. D1-D3.

[3] Casey, S. 1998. *Set Phasers on Stun and Other True Tales of Design, Technology, and Human Err*or (2nd edition), Santa Barbara, CA: Aegean Publishing Co.

[4] Bruemmer, D.J., Marble, J.L., & Dudenhoeffer, D.D. “Mutual Initiative in human-machine teams,” in *Proceedings of the IEEE Conference on Human Factors and Power Plants*, Scottsdale, AZ, September 2002.

[5] Yen, J., Yin, J., Ioerger, T., Miller, M., Xu, D., & Volz, R. “CAST: Collaborative agents for simulating teamwork,” in *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, Seattle, WA, 2001.

[6] Goodrich, M.A., & Boer, E.R. 2000. “Designing human-centered automation: Tradeoffs in collision avoidance system design,” *IEEE Transactions on Intelligent Transportation Systems, 1*.

[7] Perzanowski, D., Schultz, A.C., Adams, W., Skubic, M., Abramson, M., Bugajska, M., Marsh, E., Trafton, J. G., & Brock, D. “Communicating with teams of cooperative robots,” in *Proceedings from the 2002 NRL Workshop on Multi-Robot Systems,* Washington, D. C., March 2002.

[8] Kortenkamp, D., Huber, E., & Bonasso, R.P. “Recognizing and interpreting gestures on a mobile robot,” in *Proceedings of AAAI 1996*. Portland, OR, August 1996.

[9] Voyles, R., & Khosla, P. 1995. “Tactile gestures for human-robot interaction,” in *Proc. of IEEE/RSJ Conf. On Intelligent Robots and Systems, 3*, 7-13.

[10] Fong, T., Thorpe, C., & Baur, C. “Robot as partner: Vehicle teleoperation with collaborative control,” in *Proceedings from the 2002 NRL Workshop on Multi-Robot Systems*, Washington, D. C., March 2002.

[11] Murphy, R. R., & Rogers, E. 1996. "Cooperative assistance for remote robot supervision." *Presence, 5,* 224-240.

[12] Goodrich, M.A. & Olsen, D.R., Crandall, J.W., & Palmer, T.J. "Experiments in adjustable autonomy." in *Proceedings of the IJCAI 2001 Workshop on Autonomy Delegation and Contro*l, Seattle WA, August 2001.

[13] Scholtz, J. "Human-robot interactions: Creating synergistic cyber forces," In *Proceedings from the 2002 NRL Workshop on Multi-Robot Systems*, Washington, D. C., March 2002.

[14] Yamauchi, B., Schultz, S., & Adams, W. (1999). "Integrating exploration and localization for mobile robots," *Adaptive Behavior, 7*, 217-229.

VI. Appendix A: Questionnaire presented after each session

Please indicate the degree that you feel the following statements are true or false with reference to the last search session that you performed. Indicate your answer by circling the response the best reflects your view. Please do not mark between the numbers.

Note: "Targets" refers to the primary stuffed animal locations that you were to navigate to, while "secondary targets" refers to the extra stuffed animals that you may have seen.

1. It was very easy to send the robot through doorways and into small spaces.

Very true				Very false
1	2	3	4	5

2. I always felt that I knew where the robot was in the building.

Very true				Very false
1	2	3	4	5

3. I thought this session was fun.

Very true				Very false
1	2	3	4	5

4. It was frustrating to have to move the robot into small spaces or doorways.

Very true				Very false
1	2	3	4	5

5. I always felt that I was in control of the robot.

Very true				Very false
1	2	3	4	5

6. The amount of information presented in the interface was overwhelming.

Very true				Very false
1	2	3	4	5

7. **I used the map a lot while trying to find the targets.**
Very true Very false
1 2 3 4 5

8. **It was difficult to see objects unless the robot was pointed right at them.**
Very true Very false
1 2 3 4 5

9. **I was frequently not sure what direction the robot was traveling.**
Very true Very false
1 2 3 4 5

10. **It was difficult to predict what the robot would do when I touched the controls.**
Very true Very false
1 2 3 4 5

11. **It was difficult to tell when there were obstacles around the robot.**
Very true Very false
1 2 3 4 5

12. **I thought the video feed was very useful.**
Very true Very false
1 2 3 4 5

13. **I got frustrated using the robot at many points.**
Very true Very false
1 2 3 4 5

14. **The map did not help me to find the targets.**
Very true Very false
1 2 3 4 5

15. **I am sure that I saw all the secondary targets.**
Very true Very false
1 2 3 4 5

16. **I seldom looked at the video feed when I was navigating.**
Very true Very false
1 2 3 4 5

17. **The information I needed to control the robot to get was present and always easy to find.**
Very true Very false
1 2 3 4 5

18. **It was easy to predict how the controls would work.**
Very true Very false
1 2 3 4 5

19. **I could not always tell where obstacles were relative to the robot.**
Very true Very false
1 2 3 4 5